Detection of Sensor Irregularities in Fitness Time Series

Rayappa David Amar Raj[1], Bota Dusenbi[1], Kanasottu Anil Naik[2] and Sreenu Sunkaraboina[3]

[1]National Institute of Technology, Warangal, Telangana, India
[2]National Institute of Technology, Warangal, Telangana, India
[3]Vardhaman College of Engineering, Hyderabad, India

Abstract
Wearable fitness devices are widely used to monitor physiological signals such as heart rate (BPM) and speed during physical activity. However, these signals often suffer from noise, technical inaccuracies, and context-dependent variability. In this study, we investigate unsupervised anomaly detection methods to identify abnormal segments in real-world data collected from runners using wearable sensors. The dataset includes over 180,000 measurements from 43 running sessions, with speed and BPM values aligned and preprocessed to build a multivariate time series. We compare four approaches representative of different anomaly detection paradigms: distance-based (k-Nearest Neighbors), classification-based (One-Class SVM), probabilistic (Kernel Density Estimation), and sequence-based deep learning (TadGAN). Classical methods operate on point-wise values and capture global anomalies with high precision, but they fail to detect contextual or collective anomalies. TadGAN, in contrast, is trained on overlapping sequences and demonstrates the ability to identify local patterns of abnormality across time. Our results highlight the complementarity of these methods and the importance of modeling temporal structure when anomalies are subtle or context-dependent. Although TadGAN fails to capture extreme point anomalies, its performance on sequence-level detection suggests promising directions for future research in health-aware fitness monitoring. All analyses were conducted without labels, under purely unsupervised conditions.

Keywords
Anomaly Detection, Time Series Analysis, Wearable Sensors, Unsupervised Learning

1. Introduction

Wearable fitness technologies have become a cornerstone of modern personal health monitoring, offering non-invasive and continuous access to physiological and kinematic data during exercise. Among the most frequently collected variables are heart rate (BPM) and locomotor speed, which are commonly used as indicators of training intensity, cardiovascular load, and overall physical condition. These metrics are especially critical in endurance disciplines such as running, where performance, safety, and adaptation must be balanced in real time.

However, despite their utility, such measurements are inherently subject to various sources of uncertainty [1, 2]. These include instrumental inaccuracies (e.g., sensor resolution, sampling delay), environmental noise (e.g., GPS instability, skin reflectance), and inter-individual variability (e.g., age, fitness level, recovery status). Consequently, the raw data streams acquired from commercial devices may contain errors, inconsistencies, or even misleading information. From a data science perspective, such irregularities can be interpreted as *anomalies* — values or patterns that deviate from expected physiological behavior.

Detecting these anomalies is of dual importance. On the one hand, it allows for early identification of sensor malfunctions or transmission artifacts, which can prevent incorrect feedback or unsafe decisions. On the other hand, it may highlight abnormal physiological events, such as arrhythmias, overexertion episodes, or disruptions due to illness or fatigue. Therefore, anomaly detection is not only a technical challenge but also a potential health safeguard.

In this work, we address the task of unsupervised anomaly detection applied to a dataset of 43 running sessions. Each session contains synchronised recordings of heart rate and speed. The goal is to identify three main categories of anomalies:

- Point anomalies, which are isolated and extreme values, such as sharp spikes due to sensor loss or motion artifacts;
- Contextual anomalies, which occur when a value is inconsistent with its immediate temporal context (e.g., a sudden drop in BPM during steady running);
- Collective anomalies, which emerge as entire segments of behavior deviating from usual multivariate dynamics.

These types are illustrated in Figure 1, adapted from the survey by Ruff et al. [3].

The practical relevance of such analysis is evident in many use cases: performance tracking, adaptive training programs, fatigue detection, and sensor validation. Furthermore, since the data are unlabeled — no ground

SYSTEM 2025: 11th Sapienza Yearly Symposium of Technology, Engineering and Mathematics. Rome, June 4-6, 2025
anilnaik205@nitw.ac.in (K. A. Naik);
sreenu1792@vardhaman.org (S. Sunkaraboina)

truth exists for "normal" versus "anomalous" episodes — only unsupervised learning techniques are applicable. These include density-based, distance-based, probabilistic, classification-driven, and generative models.

From a physiological standpoint, both BPM and speed are computed indirectly. Heart rate is commonly estimated through photoplethysmography (PPG) via optical sensors on the wrist. This technique relies on green LED light absorption by pulsating blood flow and is known to be sensitive to placement, motion, and skin characteristics [4, 5]. Speed, on the other hand, is estimated using GPS signals either via positional differentiation or Doppler shift. The latter offers greater precision, particularly at high velocities [6], but even then, studies report a 3–8% error margin depending on device settings and user motion [7].

Additionally, anomalies may not only stem from the devices but also from the athlete. For example, during recovery phases, illness, or unexpected fatigue, the physiological responses may diverge from usual patterns, producing authentic yet significant anomalies. Thus, distinguishing between device-induced and body-induced anomalies is itself a meaningful analytical question.

To explore this problem, we test and compare a range of unsupervised techniques:

- Three classical models on pointwise data: k-Nearest Neighbors (kNN), One-Class SVM (OCSVM), and Kernel Density Estimation (KDE);
- A density-based cluster model: DBSCAN;
- A probabilistic generative model: Gaussian Mixture Models (GMM);
- A modern time-series deep learning model: TadGAN, a GAN-based approach designed to reconstruct temporal patterns.

Each model captures a different perspective on anomaly structure, from spatial density and decision boundaries to statistical probability and generative reconstruction. Furthermore, we compare their outputs under both static (pointwise) and sequential (time-series) representations of the same dataset, to assess the importance of temporal information in the detection of anomalies in physiological data.

2. Related Works

Anomaly detection is a classical and widely explored domain, with roots extending back centuries in statistical analysis and more recently enriched by machine learning and deep learning approaches. As discussed in Ruff et al. [3] and Nassif et al. [8], modern anomaly detection techniques can be grouped into several families: distance-based, probabilistic, classification-based, and reconstruction-based models. These approaches vary in

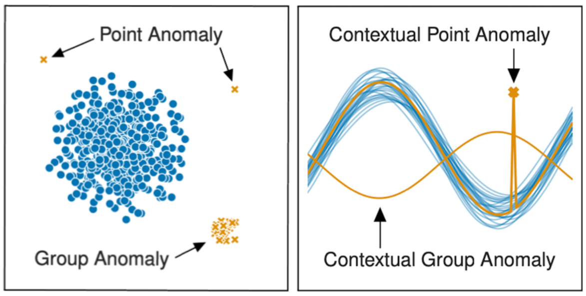

Figure 1: Illustration of anomaly types: point, contextual, and collective [3].

theoretical grounding, computational efficiency, scalability, and applicability to temporal or multivariate data.

In a broad review by Nassif et al. [9], which examines 290 research articles published between 2000 and 2020, the authors classify methods into five main categories: classification (e.g., SVM, Bayesian networks, decision trees, neural networks, kNN), clustering (e.g., k-means, hierarchical clustering), optimization-based approaches, ensemble techniques, and regression models. Often, modern systems leverage hybrid combinations of these categories to enhance robustness.

Ruff et al. [3] emphasize the theoretical motivation for unsupervised anomaly detection, focusing on the modeling of normality. This perspective leads to a division of techniques into three primary classes: probabilistic models, classification-based models, and reconstruction-based models, with distance-based methods often treated separately.

Probabilistic models attempt to fit the probability distribution of normal data and classify low-probability regions as anomalous. Classic approaches use Mahalanobis distance, Gaussian mixture models [10], kernel density estimation [11], and histogram estimators [12]. Generative models such as Variational Autoencoders (VAEs) [13] and Generative Adversarial Networks (GANs) [14] represent more recent extensions, though their accuracy often deteriorates in high-dimensional spaces unless adequately trained.

Classification models, on the other hand, explicitly attempt to separate normal from anomalous samples. One-Class SVM (OC-SVM) [15], Support Vector Data Description (SVDD) [16], and their neural extensions such as Deep SVDD [17] are well-established approaches in this family.

Reconstruction-based models rely on the premise that normal samples can be accurately reconstructed by a learned function, typically involving dimensionality reduction and encoding-decoding schemes. Common techniques include PCA [18], autoencoders, and GAN-based reconstruction [19, 20, 21, 22, 23]. Higher reconstruction errors suggest higher likelihood of anomaly, under the assumption that anomalies are rare and underrepresented in training data.

Distance-based methods, such as k-Nearest Neighbors (kNN), Local Outlier Factor (LOF), and related algorithms, operate by measuring the relative distance or density deviation of samples with respect to their neighbors. Goldstein et al. [24] showed that nearest-neighbor methods tend to outperform clustering-based methods across diverse benchmark datasets. However, they require high computational time and may not scale well to large datasets.

In temporal anomaly detection, time-series models become essential. Unlike tabular data, time-series retains sequential information, and anomalies may occur in patterns rather than isolated points. The concept of contextual and collective anomalies, discussed in Chandola et al. [25] and Al-Qassou et al. [26], becomes central. A collective anomaly corresponds to a subsequence that, while locally consistent, is globally deviant.

GAN-based models like TadGAN [27] reconstruct time series using adversarial training. TadGAN learns to map time windows to a latent representation, which is then used for reconstruction. Discrepancies between real and reconstructed sequences provide an anomaly score. This method is particularly suitable for multivariate time series, even with low dimensionality, such as our case of two correlated variables: heart rate and speed.

It is important to note that for many applications, ground-truth labels are missing, so unsupervised learning becomes necessary. Many real-world datasets require domain expert labeling, which is expensive or infeasible at scale. Therefore, unsupervised models dominate the field, often evaluated on synthetic datasets or via indirect proxy metrics.

Our specific application poses several data quality challenges. Temporal gaps, disjoint session timestamps, and asynchronous recording of heart rate and speed introduce artifacts that could mislead anomaly detection models. For example, some training sessions show BPM values while no speed is recorded, indicating different sensor systems. In other sessions, high BPM spikes do not correspond to any physical acceleration, but rather to recovery states or sensor noise.

One example is the session on 10/02, where a runner resumed activity after a one-month break. While the speed returned to prior levels, BPM values were significantly elevated. Although physiologically plausible, a model not incorporating temporal and contextual information might flag this as an error.

Thus, several methods were examined prior to implementation: - Statistical rules, such as the 3-sigma Gaussian rule. - Exploratory data techniques, such as box plots and rank analysis. - Density-based clustering, e.g., DBSCAN. - Ensemble frameworks like PyOD.

In conclusion, anomaly detection in physiological and kinetic data remains a complex, interdisciplinary task. The selected methods must balance sensitivity to real anomalies with robustness to noise. Given the nature of our dataset—low-dimensional, sequential, and partially corrupted—we selected kNN, KDE, OC-SVM for point-wise analysis, and TadGAN for sequential modeling. This hybrid approach aims to leverage both statistical precision and temporal coherence.

3. Dataset

Our dataset comprises 43 running sessions, for a total of 180,876 measurements, split into 137,515 training samples and 43,263 for testing. The data includes heart rate (beats per minute, BPM) and speed values, recorded asynchronously by different sensors and later aligned over a unified timestamp index. BPM files consist of three columns: “value”, “data”, and “startTime”. Speed files contain four columns: “data”, “startTime”, “delta_T”, and “value”. After excluding infinite values (replaced using large constants derived from the median) and applying linear interpolation to handle missing values, both signals were resampled and aggregated.

Overall, the BPM dataset contains 36,857 valid samples, while the speed dataset contains 61,748. Because the two series have different lengths and inconsistent temporal sampling (e.g., 99.2% of speed delta_T values are 2 or 3 seconds, but some anomalous values exist), a matching step was required. To fuse the datasets, we took the BPM timestamps as reference and aggregated the speed values over those intervals using a mean operation. This resulted in the loss of approximately 21% of BPM rows due to misalignment or missing speed data at corresponding timestamps. The final aligned dataset consisted of 29,147 samples.

An exploratory data analysis was conducted on both signals. The BPM values appeared within plausible physiological ranges (60–200 bpm), whereas speed values exhibited substantial variability. Outliers were initially visualized using boxplots (Figure 4), but not removed to preserve diversity and noise for anomaly detection. Speed inconsistencies and their dependence on delta_T were analyzed (Figure 3).

Principal Component Analysis (PCA) was employed to inspect the structure of the dataset and observe potential separability among patterns or noise (Figure 5). It also allowed us to monitor the effects of outlier cleaning, as shown in Figure 8.

Finally, to prepare the dataset for modeling using GANs, we applied a sliding window approach with a window length of 100 and a step size of 5. This generated overlapping multivariate sequences from the final dataset, each containing normalized BPM and speed. These se-

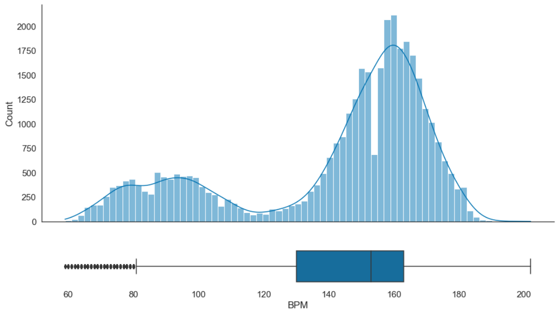

Figure 2: View on BPM.

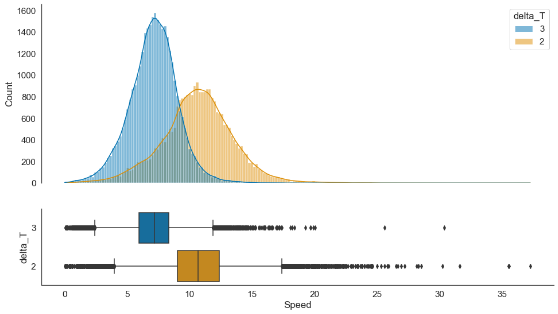

Figure 3: Repartition of the speed value depending on delta_T.

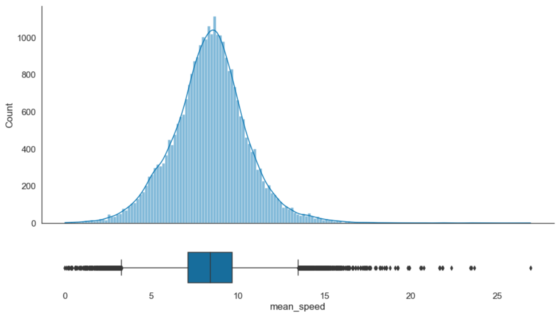

Figure 4: Outliers of speed

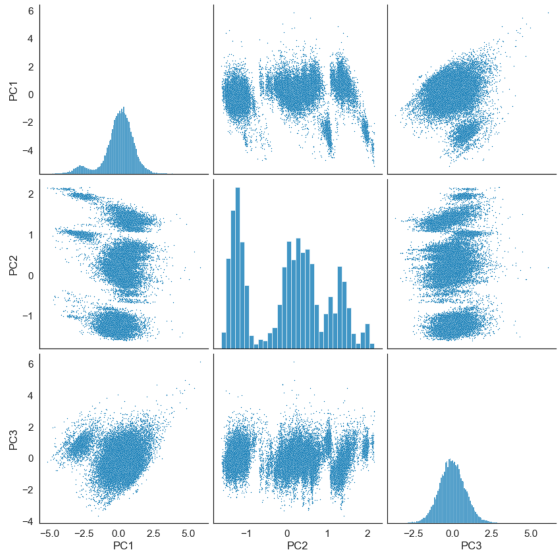

Figure 5: PCA pair plot.

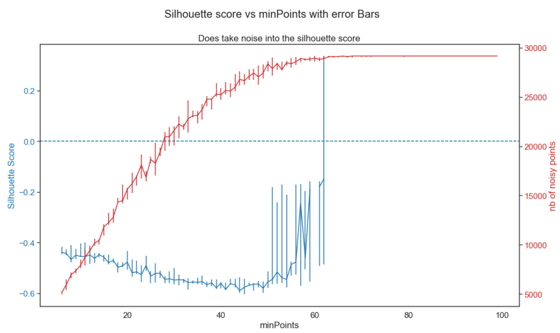

Figure 6: Silhouette score vs minPoints with error bars.

quences were used to train TadGAN.

4. Dataset

Our dataset comprises 43 running sessions, for a total of 180,876 measurements, split into 137,515 training samples and 43,263 for testing. The data includes heart rate (beats per minute, BPM) and speed values, recorded asynchronously by different sensors and later aligned over a unified timestamp index. BPM files consist of three columns: "value", "data", and "startTime". Speed files contain four columns: "data", "startTime", "delta_T", and "value". After excluding infinite values (replaced using large constants derived from the median) and applying linear interpolation to handle missing values, both signals were resampled and aggregated.

Overall, the BPM dataset contains 36,857 valid samples, while the speed dataset contains 61,748. Because the two series have different lengths and inconsistent temporal sampling (e.g., 99.2% of speed delta_T values are 2 or 3 seconds, but some anomalous values exist), a matching step was required. To fuse the datasets, we took the BPM timestamps as reference and aggregated the speed values over those intervals using a mean operation. This resulted in the loss of approximately 21% of BPM rows due to misalignment or missing speed data at corresponding timestamps. The final aligned dataset consisted of 29,147 samples.

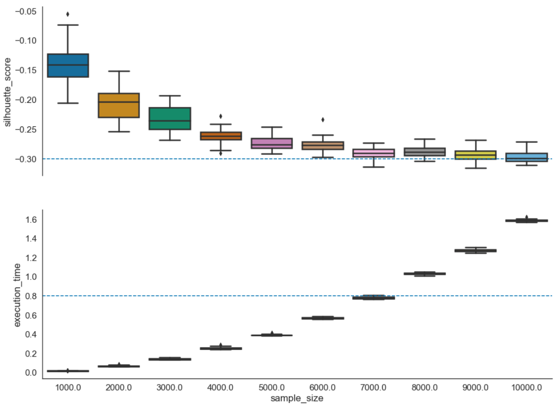

Figure 7: Trade-off between execution time and silhouette approximation.

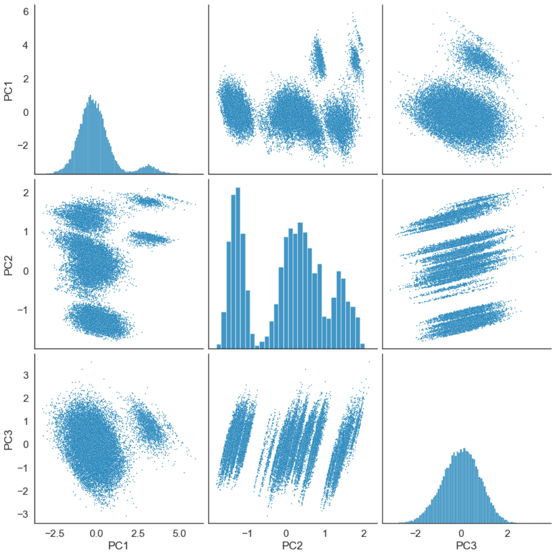

Figure 8: PCA plot after data cleaning.

An exploratory data analysis was conducted on both signals. The BPM values appeared within plausible physiological ranges (60–200 bpm), whereas speed values exhibited substantial variability. Outliers were initially visualized using boxplots (Figure 4), but not removed to preserve diversity and noise for anomaly detection. Speed inconsistencies and their dependence on delta_T were analyzed (Figure 3).

Principal Component Analysis (PCA) was employed to inspect the structure of the dataset and observe potential separability among patterns or noise (Figure 5). It also allowed us to monitor the effects of outlier cleaning, as shown in Figure 8).

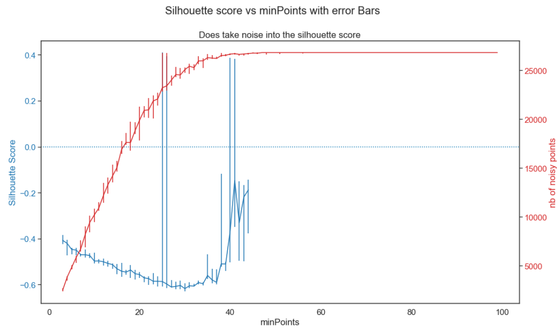

Figure 9: Silhouette score and estimated time.

Table 1
Data set description

	Train	Test	Total
# runs	34	9	43
# points	137515	43263	180876
# inf speed points	1	1	2
# inf bpm points	0	0	0
# NaN speed points	3218	101	3319
# NaN bpm points	108993	35026	144019

Finally, to prepare the dataset for modeling using GANs, we applied a sliding window approach with a window length of 100 and a step size of 5. This generated overlapping multivariate sequences from the final dataset, each containing normalized BPM and speed. These sequences were used to train TadGAN.

5. Results and Discussion

Since no ground-truth labels are available in the dataset, the evaluation of the anomaly detection methods was carried out through qualitative visual inspection of the results. In Table 2, we provide a structured comparison of anomaly scores and detection outputs from all tested methods. Each method produces scores on different numerical scales; therefore, a MinMaxScaler was applied to normalize scores in the [0, 1] interval, enabling a direct comparison.

In the first row of Table 2, we observe the distribution of the normalized anomaly scores for each method, with the red horizontal line indicating the 99.9th percentile. This threshold was then employed to classify outliers. The second row shows the test dataset over time, where speed and BPM are plotted with marked anomalies based on the threshold. The third row presents a scatter plot of BPM vs speed, where colors reflect the anomaly score, helping to distinguish point and group anomalies. Finally, the fourth row illustrates one representative test

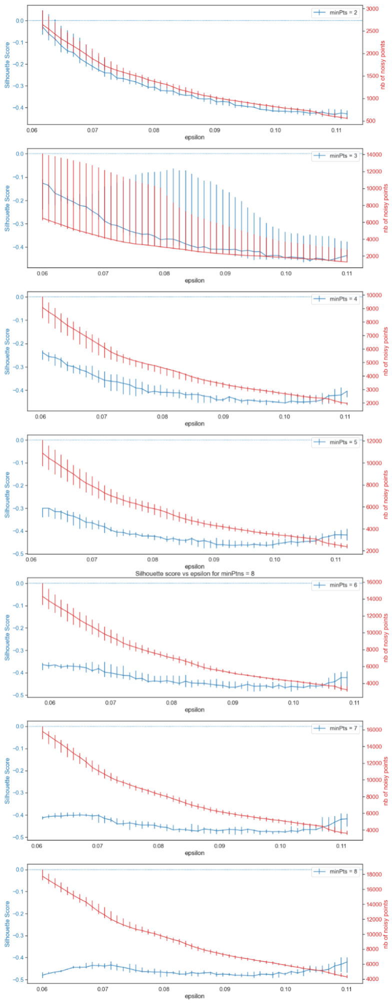

Figure 10: Silhouette score and number of Noisy points across 100 values of eps for each minPts in [2, 8].

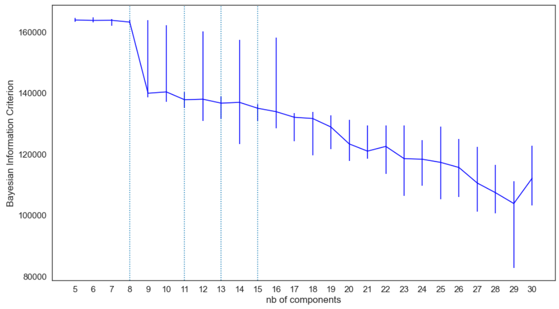

Figure 11: BIC vs number of components.

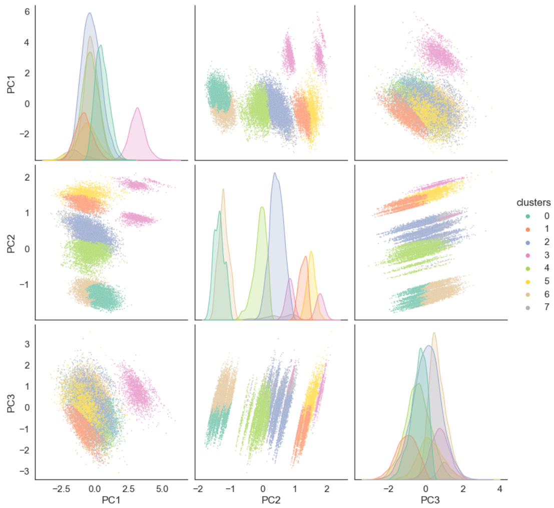

Figure 12: PCA pair plot with cluster labelling.

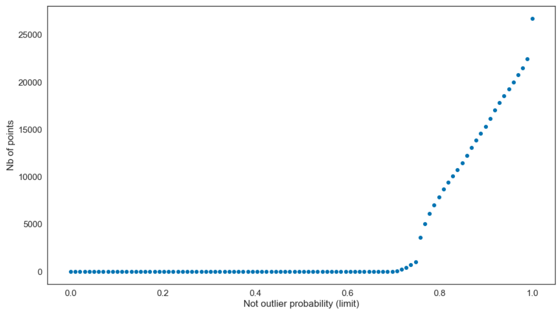

Figure 13: Number of outliers according to a tolerance level.

session in time-series format, comparing raw signals and detected anomalies.

K-Nearest Neighbors (KNN) and One-Class SVM (OCSVM) showed good capability to highlight global point anomalies but failed to capture more subtle temporal patterns or contextual outliers. Their scores remained mostly stable across time, and rapid local variations in BPM or speed were not reflected in the outputs. As a result, the methods could only detect extreme values that deviate globally from the norm.

The Kernel Density Estimator (KDE), instead, demonstrated higher sensitivity to group anomalies and moderately abnormal sequences. It successfully identified both global point outliers and sustained deviations in BPM or speed. However, due to its lack of temporal context integration, it missed several subtle dips and transient events that would qualify as contextual anomalies.

TadGAN, a generative model trained over temporal windows, exhibited a distinct behavior. The anomaly score was highly sensitive to the shape and structure of

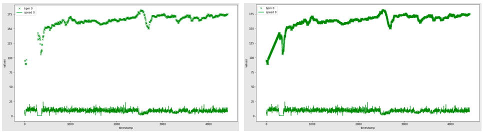

Figure 14: Before (left) and After (right) preprocessing

the sequence, favoring low variation (plateau) segments. As a result, sudden changes or transitions received lower scores, which contrasts with expectations. This behavior likely stems from the reconstruction-based scoring mechanism and the use of overlapping sliding windows. Post-processing of overlapping scores into a single per-point score may have diluted signal peaks. Further investigation is needed to refine this post-processing, possibly by reducing the sliding window step size or aggregating scores using weighted schemes.

In conclusion, point-wise methods (KNN, OCSVM, KDE) are appropriate for detecting global anomalies and extreme values. Temporal methods like TadGAN are more suitable for discovering contextually anomalous patterns, although they require more careful calibration. The integration of hybrid methods or ensembles could offer a more balanced performance across different anomaly types.

6. Conclusion

In this study, we analyzed runner data including heart rate (BPM) and speed, with the objective of identifying anomalous patterns that could indicate sensor faults, physiological irregularities, or performance deviations. We evaluated four unsupervised methods: KNN, OCSVM, KDE, and TadGAN. The first three operate on individual datapoints, while TadGAN uses sliding temporal windows.

The results showed that point-wise methods excel at detecting isolated outliers but are limited in capturing sequential patterns. In contrast, TadGAN was more sensitive to contextual anomalies and group deviations over time but struggled with extreme values. Its limitations are likely tied to the aggregation of reconstruction errors across windows.

Future work should investigate improvements in GAN-based models by tuning window size, overlap, and scoring methods. Moreover, the integration of contextual metadata or ensemble learning approaches may enhance robustness. Overall, temporal models show promising potential for enhancing anomaly detection in physiological monitoring during exercise.

Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT, Grammarly in order to: Grammar and spelling check, Paraphrase and reword. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

[1] G. Zimatore, C. Serantoni, M. C. Gallotta, L. Guidetti, G. Maulucci, M. De Spirito, Automatic detection of aerobic threshold through recurrence quantification analysis of heart rate time series, International Journal of Environmental Research and Public Health 20 (2023). doi:10.3390/ijerph20031998.

[2] M. C. Gallotta, V. Bonavolontà, G. Zimatore, S. Iazzoni, L. Guidetti, C. Baldari, Effects of open (racket) and closed (running) skill sports practice on children's attentional performance, Open Sports Sciences Journal 13 (2020) 105 – 113. doi:10.2174/1875399X02013010105.

[3] L. Ruff, R. A. Kauffmann, R. A. Vandermeulen, G. Montavon, W. Samek, M. Kloft, T. G. Dietterich, K.-R. Müller, A unifying review of deep and shallow anomaly detection, Proceedings of the IEEE (2021).

[4] J. Montes, Wearable Technology Devices: Heart Rate and Step Count Analysis, Ph.D. thesis, University of Nevada, Las Vegas, 2019.

[5] J. W. Navalta, J. Montes, N. G. Bodell, R. W. Salatto, J. W. Manning, M. DeBeliso, Concurrent heart rate validity of wearable technology devices during trail running, Plos one 15 (2020) e0238569.

[6] J. Zhang, K. Zhang, R. Grenfell, R. Deakin, Short note: On the relativistic doppler effect for precise velocity determination using gps, Journal of Geodesy 80 (2006) 104–110.

[7] M. Cardinale, M. C. Varley, Wearable training-monitoring technology: applications, challenges, and opportunities, International journal of sports physiology and performance 12 (2017) S2–55.

[8] A. B. Nassif, M. A. Talib, Q. Nasir, F. M. Dakalbab, Machine learning for anomaly detection: A systematic review, IEEE Access (2021).

[9] A. B. Nassif, M. A. Talib, Q. Nasir, F. M. Dakalbab, Machine learning for anomaly detection: A systematic review, IEEE Access (2021).

[10] L. Xiong, B. Póczos, J. Schneider, A. Connolly, J. VanderPlas, Hierarchical probabilistic models for group anomaly detection, in: Proceedings of the fourteenth international conference on artificial intelli-

KNN	OCSVM	KDE	TadGAN

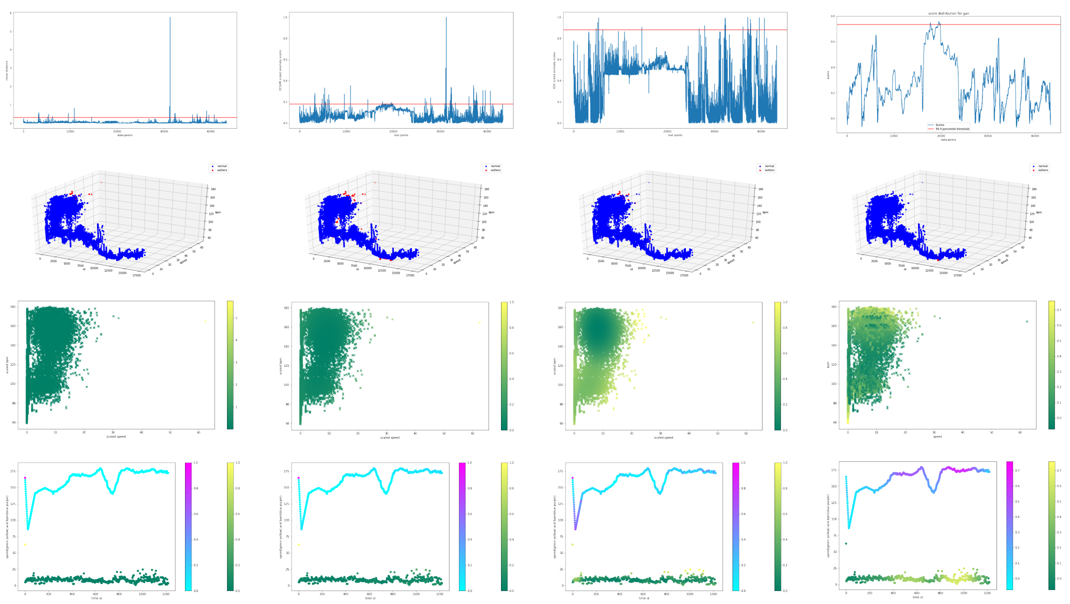

Table 2
Resultant plots of speed, BPM and anomaly scores from each method. Row-wise: (1) score distribution and 99.9th percentile threshold; (2) raw signals with detected outliers; (3) BPM vs speed with anomaly score color scale; (4) example run showing temporal anomaly scores.

gence and statistics, JMLR Workshop and Conference Proceedings, 2011, pp. 789–797.

[11] J. Kim, C. D. Scott, Robust kernel density estimation, The Journal of Machine Learning Research 13 (2012) 2529–2565.

[12] B. Nachman, D. Shih, Anomaly detection with density estimation, Physical Review D 101 (2020) 075042.

[13] J. An, S. Cho, Variational autoencoder based anomaly detection using reconstruction probability, Special Lecture on IE 2 (2015) 1–18.

[14] T. Schlegl, P. Seeböck, S. M. Waldstein, G. Langs, U. Schmidt-Erfurth, f-anogan: Fast unsupervised anomaly detection with generative adversarial networks, Medical image analysis 54 (2019) 30–44.

[15] I. Steinwart, D. Hush, C. Scovel, A classification framework for anomaly detection., Journal of Machine Learning Research 6 (2005).

[16] M. Turkoz, S. Kim, Y. Son, M. K. Jeong, E. A. Elsayed, Generalized support vector data description for anomaly detection, Pattern Recognition 100 (2020) 107119.

[17] Z. Zhang, X. Deng, Anomaly detection using improved deep svdd model with data structure preservation, Pattern Recognition Letters 148 (2021) 1–6.

[18] H. Ringberg, A. Soule, J. Rexford, C. Diot, Sensitivity of pca for traffic anomaly detection, in: Proceedings of the 2007 ACM SIGMETRICS international conference on Measurement and modeling of computer systems, 2007, pp. 109–120.

[19] H. Zenati, C. S. Foo, B. Lecouat, G. Manek, V. R. Chandrasekhar, Efficient gan-based anomaly detection, arXiv preprint arXiv:1802.06222 (2018).

[20] F. Fiani, S. Russo, C. Napoli, A fully automatic visual attention estimation support system for a safer driving experience, in: CEUR Workshop Proceedings, volume 3695, 2023, p. 40 – 50.

[21] D. Li, D. Chen, B. Jin, L. Shi, J. Goh, S.-K. Ng, Mad-gan: Multivariate anomaly detection for time series data with generative adversarial networks, in: International Conference on Artificial Neural Networks, Springer, 2019, pp. 703–716.

[22] S. Russo, I. E. Tibermacine, A. Tibermacine, D. Chebana, A. Nahili, J. Starczewscki, C. Napoli, Analyzing eeg patterns in young adults exposed to different acrophobia levels: a vr study, Frontiers

in Human Neuroscience 18 (2024). doi:10.3389/fnhum.2024.1348154.

[23] G. De Magistris, L. Guercio, F. Starna, S. Russo, N. Kryvinska, C. Napoli, A real-time support with haptic feedback for safer driving using monocular camera, in: Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), volume 15450 LNAI, 2025, p. 161 – 174. doi:10.1007/978-3-031-80607-0_13.

[24] M. Goldstein, S. Uchida, A comparative evaluation of unsupervised anomaly detection algorithms for multivariate data, PloS one 11 (2016) e0152173.

[25] V. Chandola, A. Banerjee, V. Kumar, Anomaly detection: A survey, ACM computing surveys (CSUR) 41 (2009) 1–58.

[26] R. Al-amri, R. K. Murugesan, M. Man, A. F. Abdulateef, M. A. Al-Sharafi, A. A. Alkahtani, A review of machine learning and deep learning techniques for anomaly detection in iot data, Applied Sciences 11 (2021) 5320.

[27] A. Geiger, D. Liu, S. Alnegheimish, A. Cuesta-Infante, K. Veeramachaneni, Tadgan: Time series anomaly detection using generative adversarial networks, in: 2020 IEEE International Conference on Big Data (Big Data), IEEE, 2020, pp. 33–43.